DOCUMENT RESUME

ED 462 444 TM 033 709

AUTHOR Anderson, John O.
TITLE The Evaluation of Student Achievement: Preliminary Analyses in Modelling Teacher Decisions.
PUB DATE 2000-05-00
NOTE 15p.; Paper presented at the Annual Meeting of the Canadian Society for the Study of Education (Alberta, Canada, May 24-28, 2000.
PUB TYPE Reports - Research (143) -- Speeches/Meeting Papers (150)
EDRS PRICE MF01/PC01 Plus Postage.
DESCRIPTORS *Academic Achievement; Decision Making; Elementary Education; Elementary School Students; *Elementary School Teachers; Grades (Scholastic); *Grading; *Preservice Teachers; Student Evaluation; Student Journals; Teacher Attitudes; Teacher Education

ABSTRACT

This study focused on one task that is characteristic of teacher responsibilities and activities in the school: the evaluation of student achievement. It involved more than 100 preservice elementary school teachers who assessed the performance of 3 simulated students on 6 language arts tasks. Information collected included the marks assigned to students on various submitted assignments and tests and the journal entries of the student teachers. The study continues an investigation into the procedures and information bases preservice teachers use in making judgments about student achievement. The marks and grades that each student teacher generated were summarized and compared across the three simulated students to determine the extent to which the student teachers viewed their three students as distinct in their achievement in language arts. Study findings support the view that evaluation of student achievement is not a simple process. The data show that final marks are not the same thing as final letter grades, although they are closely related. Educators have characteristic predilections to mark or grade high or low, and elements other than marks awarded to specific achievement products enter into the creation of final marks and letter grades. Results also demonstrate the potential of the portfolio approach to collecting information about the evaluation of student achievement by teachers. The achievement products in this study appear to have functioned as intended, in that the expected student achievement level was recognized by the evaluators. Analysis of the journals of the student teachers will be used to provide further information about how teachers evaluate students. (Contains 7 tables and 22 references.) (SLD)

Reproductions supplied by EDRS are the best that can be made
from the original document.

ED 462 444

The Evaluation of Student Achievement: Preliminary Analyses in Modelling Teacher Decisions

John O. Anderson

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

X This document has been reproduced as received from the person or organization originating it.

☐ Minor changes have been made to improve reproduction quality.

• Points of view or opinions stated in this document do not necessarily represent official OERI position or policy.

PERMISSION TO REPRODUCE AND DISSEMINATE THIS MATERIAL HAS BEEN GRANTED BY

J. Anderson

TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

TM033709

BEST COPY AVAILABLE

The Evaluation of Student Achievement:
Preliminary Analyses in Modelling Teacher Decisions

John O. Anderson
University of Victoria

What teachers know, how they know it and how this relates to practice seems, on the surface, to be an area readily accessible to investigation and the answers should be largely in place. However, this is not the case. The study of teacher knowledge is plagued by at least two fundamental problems: proximity and situational complexity. To an extent most people know something about teacher knowledge, having been exposed to teachers and teaching for over a dozen years of formal schooling. Most readers of papers such as this (who would be university researchers) have a closer association with teacher knowledge in that they have a store of their own teacher knowledge, having been an educators most of their adult (post-schooling) life. This proximity can be viewed as a fundamental problem in investigating any aspect of education and teacher knowledge in particular. University researchers may be simply too close to the phenomena to clearly view it in others and if we do consider teacher knowledge as a distinct entity we are likely to assume that we more or less know what it is and how it works. This can distort the meanings of what we do chose to see and hear. A compounding problem is the nature of teacher (or any human) knowledge – it is dynamic, complex and situationally specific. In other words the phenomenon of teacher knowledge has not generally been well studied. However, over the past decade or more there has been considerable work completed on researching teacher knowledge as evidenced in major reviews such as those of Grimmett and MacKinnon (1992) and Fenstermacher (1994). It appears that teacher knowledge is generally studied through narrative and reflective case study approaches in which teacher knowledge is investigated as a whole. This contrasts to earlier research approaches (Dunkin & Biddle, 1974; Brophy & Good, 1984) that were of a more positivistic character and more analytic in their decomposition of knowledge of teaching and teacher knowledge into constituent parts (variables). Many of the more recent studies are conducted within the context of reform of teacher education (Grimmett, 1998; Hargreaves & Jacka, 1995) or schools in general (Clandinin & Connelly, 1987; Wideen, Mayer-Smith & Moon, 1996). As noted by Carter (1993), the situational specificity of cases and the multiple interpretations of stories can lead to problems in developing a corpus of research that will cogently inform teacher education or school reform. So, there is much work yet to be done in the area of teacher knowledge and the study reported here is part of a collaborative attempt to make a positive contribution to understanding teaching and learning.

The collaborative research program has been reported elsewhere (Wilson, 1999) and in several papers presented at the 2000 *Canadian Society for the Study of Education* conference (Shulha, 2000; Locke, 2000; Wilson, 2000; Petrick, 2000; Notman, 2000; Lee, 2000; Muir, 2000). This program has narrowed the focus of investigation to classroom assessment practices yet maintained a rather broad scope of investigatory approaches – including case study participant research, journal based narrative analysis, and both inferential and descriptive modelling statistical analyses – within a collaborative context (Shulha, Wilson & Anderson, 1999).

The Study

The study reported in this paper isolated its focus on one task that is characteristic of teacher responsibilities and activities in the school: the evaluation of student achievement. It involved over one hundred pre-service elementary teachers who assessed the performance of three simulated students on a number of language arts tasks. Information collected included the marks assigned to students on various submitted assignments and test, and journal entries.

The study continues an investigation into the procedures and information bases pre-service teachers use in making judgements about student achievement (Wilson & Martinussen, 1999; Shulha, 1999; *Author*, 1999). The initial studies attempted to investigate how teachers go about evaluating the achievement of their students. These studies were based upon a dataset developed by Wilson and Shulha of Queen's University. They created a set of portfolios containing achievement products (such as written assignments and tests) and background information for a simulated student called *Chris* in a grade 8 language arts class. The contents of the portfolio were controlled in terms of achievement level of products and the background of the student. This resulted in three different simulated students. Each participating student teacher was given the work of one of the simulated students to evaluate. As part of an undergraduate teacher education course in classroom assessment, 147 student teachers graded the components of an assigned portfolio over a 12-week period and reported a final grade for *Chris* at the end of the term. These scores and grades were the basis for an investigation of the structure underlying the evaluation of achievement by these student teachers.

The current study utilized a traditional empirically based research approach to analysing the scores and grades generated by the student teachers. The portfolio structure was modified in that each portfolio contained the work of three different students on six language arts tasks, and each of the more than 100 student teachers graded the same three students. The students were now assumed to be in grade 5. Although the language arts assignments and tests were essentially the same tasks as in the previous study, student responses were modified to better reflect the work of grade 5 students. Each student teacher was required to grade each assignment as if it was requested by their sponsor teacher. Accompanying each set of student responses were instructions from the simulated sponsor teacher in regard to the grading (for example, some background to the student tasks and the total worth of each assignment). However directions in regard to how to grade the student work were designed to be rather ambiguous. The student teachers were not provided with marking criteria, keys or rubrics. Student teachers were also required to maintain a journal in which they recorded the thoughts they had about the work they were doing with their portfolios. It was suggested that any comments, views, frustrations and accomplishment they encountered in marking the student work was to be noted and discussed in their journal.

The basic data layout consists of a single complex record for each participating pre-service teacher (Figure 1). Each record contains the same data elements but varied in terms of content and structure - particularly the journal entries since there was wide variation in the nature and volume of the information

written by participants. The analysis of this information involved both statistical and interpretive approaches. The study reported in this paper consisted of the analysis of the marks and grades only.

Journal Entries	**Student 1**	**marks & grades**
	Student 2	**marks & grades**
	Student 3	**marks & grades**

Figure 1: Data layout for each student teacher record

The marks and grades that each student teacher generated were summarized and compared across the three simulated students. The goal was to investigate the extent to which the student teachers viewed their three students as distinct in terms of their achievement in language arts. The design of the portfolios was intended to create a low achieving student, one that was high and another who was a mid-range achiever through the development and inclusion of student work that consistently represented what was viewed as low, mid and high ranges of achievement. The extent to which these results are reflected in the grades and marks assigned by the student teachers could be considered an index of the design representativeness of the portfolios.

The intercorrelations of marks and grades were calculated to investigate the extent to which each assignment and test yields the same kind of information about the simulated student. Since the underlying factor in the student work is language arts achievement, it was anticipated that strong, positive correlations within each student's set of marks would emerge. Factor analyses were also conducted to investigate latent structures in these data.

The journal entries have been explored and analyzed to reveal elements and patterns in the thoughts, concerns and issues that student teachers expressed as they were attempting to complete their task of grading their three students (Bachor & Baer, 2000). The student teachers were all given the same materials on the students, the cooperating teacher and the school. Since this information was rather sparse, there were likely to be variant interpretations of the task and situation. Ambiguities of expectations and task definitions were issues that were expected to be expressed in the journals. As well, the rather limited information provided on each of the three students in the portfolio has created a more decontextualized evaluation situation than what is likely to occur in most classrooms. It is expected that the extent to which this is noted as an issue in the journals may be a major element of the journal data. However, the analysis of the journal entries may provide a rich source of information about the concerns and thoughts related to the evaluation of student achievement.

Since the journal entries are linked to the marks and grades the patterns in the journal data can inform the further analysis of the marks and grades. Categorizing meaningful patterns found within the journal data will allow for the use of both teacher perspectives from the journal entries and assigned marks in statistically modeling the evaluation of student achievement which will constitute the next stage of the investigation. This will prove to be a complex task. The use of journal entries for the development of categorical information will be fed into a structural equation model, this should allow for the development of a model that is based upon structures suggested by the thinking of the individuals generating the achievement data. The previous studies in this research embedded information about the simulated student into the portfolio materials. The model developed from these data (*Author*, 1999) was meaningful but accounted for a relatively small proportion of variance in the assessment data. It is anticipated that the results of this study should provide a basis to the development of a model that will facilitate the study of the structures underlying the evaluation of student achievement.

The Instrument

Each portfolio contained the responses of three simulated students to six language arts tasks:

1. *A Trip to the Mall* – a brief essay about going to the mall that was to be handed in as a printed word processing document. Student teachers were asked to mark this out of 12 and focus their attention on written expression rather than computer competency.
2. *Did I Order an Elephant?* – A worksheet consisting of a cloze-type reading task in which students were required to generate the 15 missing words in a reading passage.
3. *A Salmon for Simon* – A worksheet that was a modified cloze reading task in which students were required to correctly select from 5 embedded multiple-choice alternatives a phrase to complete the text, followed by 4 multiple-choice comprehension items.
4. *The New Kid on the Block* – A worksheet requiring students to read a passage and then answer 6 short answer items in which the student had to interpret a phrase from one of the character's perspective and translate into their own words a quote from the story.
5. *The Mending Wall* – A writing task in which students had to read a 43-line poem and then write a piece describing the personal meaning they found in the poem after having developed a web outlining the main ideas in the poem and discussing this with their teacher. This was to be marked out of 25.
6. *Final Exam* – A formal written exam consisting of 20 word identification items (classify word as a noun, adjective, verb or adverb), a paragraph in which the student had to extract 5 nouns and 5 verbs, 14 editing items for commas and correct capitalization, and a reading passage followed by 8 multiple-choice comprehension

items and 2 written responses (5 and 10 lines of space provided for student response).

For each of these tasks, three responses or achievement products were created for inclusion in the portfolio. One product was developed to represent low achievement, another was created to represent average or moderate achievement, and one was created to represent high achievement. The development of the achievement products (simulated student responses) involved both the researchers and some elementary school students who developed and located the level of responses. The low-level products were consistently assigned to *Student A*, the mid-level products assigned to *Student C* and the high-level products assigned to *Student B*. The effect intended was that each student teacher would assemble a portfolio of achievement products for three students: a high achiever, a moderate achiever and a low achiever. The summary statistics (Table 1) and the plots (Figure 2) show that this design expectation was realized for most of the achievement products. One exception was *A Trip to the Mall*, the first product given to the

Table 1: Summary Statistics for Marks Awarded Students A, B & C.

TASK	STUDENT A Mean (SD)	B Mean (SD)	C Mean (SD)
Trip to Mall	8.0 (1.5)	9.5 (1.3)	10.9 (1.2)
Did I Order an Elephant?	5.9 (0.8)	7.6 (0.6)	6.6 (0.8)
Salmon for Simon	5.1 (1.0)	7.4 (0.8)	5.1 (1.0)
New Kid on the Block	13.2 (2.6)	16.3 (1.7)	14.4 (2.3)
The Mending Wall	12.6 (3.2)	23.2 (1.6)	20.0 (2.4)
Final Exam	22.6 (3.6)	46.2 (3.7)	37.6 (3.2)

student teachers for inclusion in the portfolio where Student C (the moderate achiever) was generally awarded higher marks than our high achiever (Student A). The other exception was the similar results for Student A (low) and Student C (moderate) on their results for *A Salmon for Simon* worksheet. There were no tasks on which the average scores of the low achiever (A) were greater than those of the high achiever (B). On the basis of these descriptive findings it was concluded that the achievement products included in the assembly of the portfolio were representative of low, moderate and high achieving students.

Figure 2: The Results for Students A, B & C on the Tasks

Figure 2a Results Trip to the Mall Task

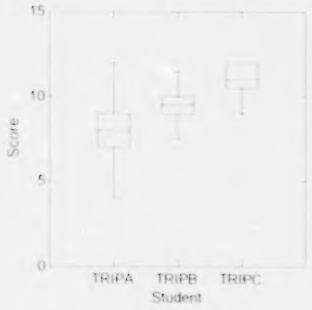

Figure 2b Results Did I Order an Elephant?Task

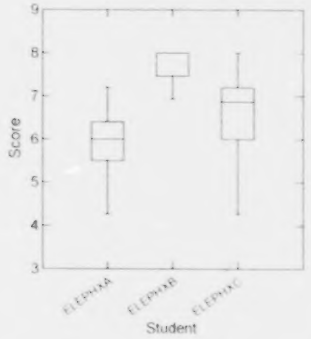

Figure 2c Results for Salmon for SimonTask

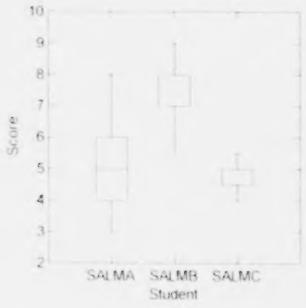

Figure 2d Results for New Kid on BlockTask

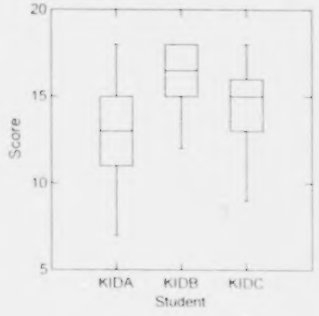

Figure 2e Results for The Mending Wall Task

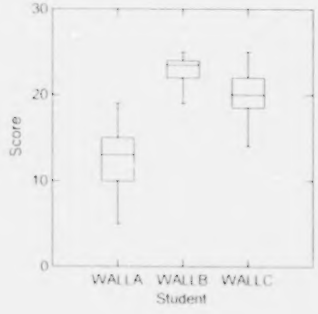

Figure 2f Results for Final Exam

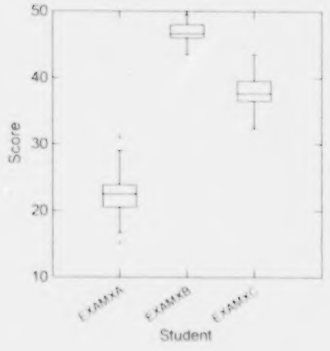

Results

Student teachers were requested to mark each of the six achievement products of each of the three students and then submit a final mark for each student. Eighty-two student teachers submitted both a *Final Mark* (a numerical score) and a *Lettergrade* (104 submitted a *Final Mark* only). For analysis purposes the *Lettergrades* were transformed into numbers with A+ being given a value of 7, an F (fail) being assigned a 1 and the rest of the grades ranged accordingly in between.

In the marking of two of the six assignments, student teachers varied a bit in the maximum scores allowed. For the task *Did I Order an Elephant?*, maximum scores ranged from 5 to 45 with most student teachers using a maximum of either 8 or 15. For the *Final Exam* maximum scores ranged from 41 to 82 with most student teachers using a maximum of 50. For analysis purposes all scores were transformed to a common scale: *Did I Order an Elephant?* was scaled to a maximum of 8, and the *Final Exam* was scaled to a maximum of 50.

In considering the summary statistics it is apparent that student teachers assigned a wide range of scores to the three students on the various tasks included in the portfolio. For any particular task there is overlap in the scores assigned students A, B and C. However the ranking of the three students who were simulated in the portfolio was consistent. A crosstabultation of the final lettergrades indicated that Student B was always rated first in achievement, Student C second and Student A was consistently ranked last. A distribution of lettergrades (Table 2) shows considerable consistency across student teacher evaluations with Student B having a modal grade of A, Student C having a mode of B and Student A a modal grade of C- - results which are consistent with the design of the portfolio contents. However there is substantial variation in the grades awarded by different student teachers – for example, five student teachers awarded the highest achieving student a grade of B, the grade awarded to the lowest achieving student by three other student teachers.

Table 2: Lettergrade Distributions

Student	Grade						
	F	C-	C	C+	B	A	A+
High (Student B)	-	-	-	-	5	73	4
Moderate (Student C)	-	1	-	8	66	6	-
Low (Student A)	10	41	23	4	3	-	-

Table 3: Correlations between achievement products for the three students

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
1 TRIPA	1.00																
2 TRIPB	0.56	1.00															
3 TRIPC	0.27	0.47	1.00														
4 SALMA	0.05	0.02 -	0.11	1.00													
5 SALMB	-0.03	0.05	-0.10	0.20	1.00												
6 SALMC	0.02	-0.09	-0.13	0.56	0.08	1.00											
7 KIDA	-0.05	0.04	-0.21	0.22	0.15	0.07	1.00										
8 KIDB	0.02	0.15	-0.06	0.21	0.12	-0.05	0.30	1.00									
9 KIDC	0.16	0.06	-0.09	0.16	0.01	0.03	0.35	0.35	1.00								
10WALLA	0.12	0.03	0.00	-0.05	0.02	-0.09	-0.02	-0.00	0.12	1.00							
11WALLB	0.10	0.04	-0.07	0.12	0.16	0.04	0.10	0.10	0.11	-0.02	1.00						
12WALLC	0.19	0.09	0.05	-0.05	-0.02	-0.03	-0.14	0.04	0.21	0.29	0.39	1.00					
13ELLPHXA	0.09	0.22	-0.01	0.18	0.05	0.04	0.02	0.06	0.05	0.05	0.14	0.06	1.00				
14ELLPHXB	0.24	0.22	0.09	0.07	0.14	0.06	0.20	0.15	0.13	-0.00	0.29	0.21	0.28	1.00			
15ELLPHXC	0.22	0.17	-0.00	0.18	0.09	0.09	0.16	0.11	0.19	-0.11	0.34	0.09	0.36	0.65	1.00		
16EXAMXA	0.06	0.08	0.15	0.05	0.26	-0.04	0.13	0.15	0.19	0.18	0.01	0.21	0.03	0.25	0.04	1.00	
17EXAMXB	0.13	0.13	0.01	0.15	-0.01	0.08	-0.01	0.00	0.01	0.00	0.15	-0.00	0.12	0.21	0.29	-0.22	1.00
18EXAMXC	0.07	0.06	0.10	0.03	-0.02	0.11	0.10	0.04	0.04	0.15	-0.15	-0.00	0.12	0.12	0.07	0.30	0.25

Correlations were calculated for all marks (Table 3) to investigate relationships between achievement tasks and among students. Preliminary perusal of the correlations suggests that there are higher correlation between instruments than between students. For example, the correlations between *Trip to the Mall* for students A, B and C are 0.56, 0.27 and 0.47 whereas the correlation between marks awarded to a given student on this assignment are all lower than 0.27. This pattern is typical of the correlation structure for these results and indicates that a high student score on one assignment did mean a high score for that student on another assignment, but a high score awarded to Student A on assignment one would be related to a high score awarded by that student teacher to Student B on that same assignment. This *marker tendency* was further revealed by the correlations of final marks for each of the three students (Table 4). The correlations are all positive ranging form 0.21 to 0.47 suggesting that student teachers who award a high final mark to Student A will also tend to award a high final mark to Students B and C, whereas a student teacher awarding a low final mark to Student A will tend to award a low final marks to Students B and C.

Table 4: Correlations of Final Marks

Student	A	B	C
A			
B	0.21		
C	0.47	0.37	

As it turned out, although there was a strong positive relationship between the final mark and the lettergrade awarded it was nowhere near perfect (Table 5). The final mark accounted for 38 to 75% of the variance in the lettergrade awarded, depending on Student. Since it was of interest to explore the nature of

Table 5: Correlation between Final Marks and Lettergrades

	Correlation
Student A	0.87
Student B	0.62
Student C	0.79

the elements contributing to the final marks and lettergrades awarded to the three students by the student teachers regression analyses were conducted using the six achievement products to predict the final mark and the lettergrade for each of the three students in the portfolio. The results (Tables 6 & 7) show that the *Final Marks* are better accounted for (R^2's range from 0.80 to 0.87) than *Lettergrades* (R^2's range from 0.31 to 0.66). The final results, particularly for lettergrades, are in some way constructed differently for each of the three students. For example, the 66% of the variance of *Lettergrades* awarded Student A (the low achieving student) are accounted for by the marks awarded the six achievement products whereas for Student B (the high achiever) only 31% of *Lettergrade* variance is accounted for by the marks and almost 70% is from other sources.

Table 6: Regressions of Marks on Achievement Products to Final Mark for Students A, B and C

	Coefficient (β)			p			R^2		
Achievement Product	A	B	C	A	B	C	A	B	C
Trip to Mall	0.20	0.17	0.15	.00	.00	.01			
Salmon for Simon	0.22	0.12	0.14	.00	.03	.01			
New Kid	0.35	0.28	0.36	.00	.00	.00			
Mending Wall	0.44	0.30	0.38	.00	.00	.00			
Elephant	0.17	0.10	0.17	.00	.07	.00			
Final Exam	0.50	0.63	0.55	.00	.00	.00			
							0.89	0.80	0.80

In summary it can be said that the *Final Mark* (the numerical final result) is well accounted for by the marks awarded on the six achievement products, and that each of the six achievement products contributed significantly to the *Final Mark*. Further, the relationships of achievement products to *Final Mark* are consistent across three students. However for the *Lettergrade*, the final results is not as well accounted for by the six achievement products. Further the extent to which achievement products contribute to the final result varies one student to another and only two of the six achievement products (*Mending Wall* and the *Final Exam*) significantly contribute to the final lettergrade. This suggests that the numerical version of the final result is not quite the same thing as the *Lettergrade* although conceptually they should convey the same information.

Table 7: Regressions of Marks on Achievement Products to Lettergrade for Students A, B and C.

Achievement Product	Coefficient (β)			p			R^2		
	A	B	C	A	B	C	A	B	C
Trip to Mall	0.22	0.02	0.04	.01	.87	.68			
Salmon for Simon	0.15	0.05	0.09	.06	.65	.35			
New Kid	0.19	-0.04	0.17	.02	.73	.09			
Mending Wall	0.41	0.34	0.34	.00	.00	.00			
Elephant	0.18	-0.03	0.18	.03	.81	.06			
Final Exam	0.33	0.34	0.40	.00	.00	.00			
							0.66	0.31	0.51

In Closing

The study supports the view that the evaluation of student achievement is not a simple process. The data shows clearly that final marks are not the same thing as final lettergrades although they are closely related. Educators have characteristic predilections to mark or grade high or low – *marker tendency* – which corresponds I am sure to many students' recollections of grades past. Further, elements other than the marks awarded to specific achievement products (worksheets, assignments and tests) enter into the creation of the final marks and lettergrades teachers assign to students, and more additional information is added into the creation of lettergrades than into numerical final marks. And finally, that the information used in the composition of lettergrades varies one student to another.

The results indicate the potential for the portfolio approach to collecting information about the evaluation of student achievement by teachers. The achievement products created for this portfolio appear to have functioned in the manner intended in that the low achieving student was perceived to be low, as did the high achieving student and the moderately achieving student.

The next steps in this research will focus on these as yet unknown information elements teachers used to develop their grades for students. To do this the information written by the student teachers in their journal should provide insightful. The patterns revealed in the analysis of the journals (Bachor & Baer, 2000) will be used to investigate the structures underlying the lettergrades assigned to our three students of varying achievement.

BEST COPY AVAILABLE

References

Author. (1999). Modeling the development of student achievement.

Bachor, D.G. & Baer, M. (2000). *An examination of pre-service teachers' portfolio diaries*. A paper presented at the annual meeting of the Canadian Society for the Study of Education, University of Alberta, May, 2000.

Brophy, J.J. & Good, T. (1986). Teacher behavior and student achievement. In M.C. Wittrock (Ed.), *Handbook on Research in Teaching (3rd edition, 328-375). New York: Macmillan.*

Carter, K. (1993). The place of story in the study of teaching and teacher education. *Educational Researcher*, 22(1), 5-12, 18.

Clandinin, D.J. & Connelly, F.M. (1996). Research in progress: studying teacher knowledge. *Alberta Journal of Educational Research*, 52(1), 67-69.

Dunkin, M.J. & Biddle, B.J. (1974). *The Study of Teaching*. New York: Holt, Rinehart & Winston.

Fenstermacher, G.D. (1994). The knower and the known: The nature of knowledge in research on teaching. In L. Darling-Hammond (Ed.) *Review of Research in Education, Vol.20*. Washington: AERA.

Grimmett, P.P. (1998). Reconceptualizing the practice of teacher preparation: on not throwing out the concurrent model with the reform bathwater. *The Alberta Journal of Educational Research*, 44(3), 251-267.

Grimmett, P.P. & MacKinnon, A.M. (1992). Craft knowledge and the education of teachers. In Grant, G. (Ed) *Review of Research in Education, Vol. 18*. Washington: AERA.

Hargreaves, A. & Jacka, N. (1995). Induction or seduction: postmodern patterns of preparing to teach. *Peabody Journal of Education*, 70(3), 41-63.

Lee, M. (2000). *The emergence and anatomy of collaboration in a teacher researcher assessment project: A cognitive and cultural view*. A paper presented at the annual meeting of the Canadian Society for the Study of Education, University of Alberta, May, 2000.

Locke, C. (2000). *Student control and responsibility: The influence of two teachers' classroom assessment*. Paper presented at the annual meeting of the Canadian Society for the Study of Education, University of Alberta, May, 2000.

Muir, W. (2000). *Collaborative research and the internet: A case study*. Paper presented at the annual meeting of the Canadian Society for the Study of Education, University of Alberta, May, 2000.

Notman, D. (2000). *The effects of portfolio assessment and student-led conferences on ownership and control*. Paper presented at the annual meeting of the Canadian Society for the Study of Education, University of Alberta, May, 2000.

Petrick, B. (2000). *The relationship between learning and assessment in the elementary classroom*. Paper presented at the annual meeting of the Canadian Society for the Study of Education, University of Alberta, May, 2000.

Shulha, L.M. (1999). Understanding novice teachers' thinking about student assessment. *Alberta Journal of Educational Research*, 45(3), 288-301.

Shulha, L.M. (2000). Implementing newer principles of assessment into planning and delivery of instruction. Paper presented at the annual meeting of the Canadian Society for the Study of Education, University of Alberta, May, 2000.

Shulha, L.M., Wilson, R.J. & Anderson, J.O. (1999). Investigating teachers' assessment practices: Exploratory, non-foundationalist, mixed method research. *Alberta Journal of Educational Research*, 45(3), 304-314.

Wideen, M.F., Mayer-Smith, J.A. &Moon, B.J. (1996). Knowledge, teacher development and change. In Ivor, F. & Hargreaves, A (Eds.) *Teachers' Professional Lives*. London: Falmer.

Wilson, R.J. & Martinussen, R.L. (1999). Factors affecting the assessment of student achievement. *Alberta Journal of Educational Research*, 45(3), 267-277.

Wilson, R.J. (1999). Classroom assessment investigations. *Alberta Journal of Educational Research*, 45(3),263-266.

Wilson, R.J. (2000). *A model of assessment-in-practice*. Paper presented at the annual meeting of the Canadian Society for the Study of Education, University of Alberta, May, 2000.

BEST COPY AVAILABLE